



Software Engineering Institute

Resource Allocation in Dynamic Environments

Jeff Hansen
Scott Hissam
B. Craig Meyers
Gabriel Moreno
Daniel M. Plakosh
Joseph Seibel
Lutz Wrage

October 2012

TECHNICAL REPORT
CMU/SEI-2012-TR-011
ESC-TR-2012-011

Research, Technology, and System Solutions Program

<http://www.sei.cmu.edu>



Copyright 2012 Carnegie Mellon University.

This material is based upon work funded and supported by under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

This report was prepared for the

SEI Administrative Agent
AFLCMC/PZE
20 Schilling Circle, Bldg 1305, 3rd floor
Hanscom AFB, MA 01731-2125

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

TM Carnegie Mellon Software Engineering Institute (stylized), Carnegie Mellon Software Engineering Institute (and design), Simplex, and the stylized hexagon are trademarks of Carnegie Mellon University.

* These restrictions do not apply to U.S. government entities.

Table of Contents

Acknowledgments	vii
Abstract	ix
1 Introduction	1
2 Background	2
2.1 Scenario Approach	2
2.1.1 Operational Configuration	2
2.1.2 Experimental Scenario	3
2.2 Resource Allocation in the Presence of Overstatement of Resource Needs	5
2.2.1 Computational Mechanism Design	6
2.3 Adaptive Quality of Service (AQoS)	8
2.3.1 QoS Resource Allocation Model (Q-RAM)	8
2.3.2 Centralized Q-RAM	9
2.3.3 Single Hop D-Q-RAM	10
2.4 Integrating the Approaches	13
3 Scenario Implementation	14
3.1 Configuration	14
3.1.1 Tactical Operations Center Configuration	15
3.1.2 Squad Configurations	16
3.1.3 Information Flows	17
3.2 TOC Camera Allocation	20
3.2.1 Eliciting Relevant Information from Squads	20
3.2.2 Computing Asset Allocation	23
3.2.3 Enacting the Allocation	24
3.2.4 AQoS as a Function of Weight	24
4 Experimental Results	27
4.1 Conducting the Experiment	27
4.2 Resource Allocation Results	27
4.2.1 Allocation Without Use of a Mechanism	27
4.2.2 Allocation that Includes a Mechanism	28
4.3 AQoS Results	30
4.3.1 TOC Camera	31
4.3.2 Ground Vehicle Camera	32
4.3.3 Squad Camera	33
4.3.4 Dropped Frames	34
4.4 Summary of Experimental Results	35
5 Summary	36
Acronym List	37
Appendix A: Scenario Details	39
Appendix B: Utility Curves Used in the Scenario	41
References	45

List of Figures

Figure 1:	Operational Configuration	2
Figure 2:	Overview of Experimental Scenario	4
Figure 3:	Markov Decision Process	7
Figure 4:	D-Q-RAM with Single Router	10
Figure 5:	Algorithm for Distributing Marginal Utility Messages on a Router	12
Figure 6:	Aerial View of Scenario Configuration	14
Figure 7:	Logical Layout of Components of Wired and Wi-Fi Network	15
Figure 8:	Screenshot from TOC Display	16
Figure 9:	Squad Interface for Requesting Resource Allocation	17
Figure 10:	Control and Resource Management Flows	18
Figure 11:	AQoS Managed Flows	19
Figure 12:	TOC Camera Allocation Sequence Diagram	20
Figure 13:	Initial MDP Template	21
Figure 14:	Simplified MDP Template	23
Figure 15:	TOC Camera Allocation Time-Averaged Utility Scores	32
Figure 16:	Ground Vehicle Camera Time-Averaged Utility Scores	33
Figure 17:	Squad Camera Time-Averaged Utility Scores	34
Figure 18:	Fraction of Frames Dropped From Left Squad Camera	35
Figure 19:	Utility Curve for the TOC Camera	42
Figure 20:	Utility Curves for Ground Vehicle Camera and Squad Camera	43
Figure 21:	Facial-Recognition Utility Values	43

List of Tables

Table 1:	Summary of Network Flows	18
Table 2:	Relative Importance of Network Messages	19
Table 3:	Parameters Used in the MDP	21
Table 4:	Mission End-State Values	22
Table 5:	Weight Sets Used in Experiment	25
Table 6:	Incentive Alignment Baseline (No Overstatement)	28
Table 7:	Squad Left Overstatement Without Mechanism for Incentive Alignment	28
Table 8:	Squad Left Overstatement with Mechanism for Incentive Alignment	29
Table 9:	No Overstatement with Mechanism for Incentive Alignment	29
Table 10:	Simulation Results of Incentive Alignment	30
Table 11:	Details of Scenario Steps	39
Table 12:	Sample Utility Values for TOC Camera	41

Acknowledgments

We acknowledge discussions with our colleagues Dionisio de Niz and Mark Klein of the SEI as well as with Michael Clement of the Naval Postgraduate School.

Abstract

This technical report examines two challenges related to resource allocation that can negatively affect system operation in a dynamic environment, where warfighter needs for resources, resource availability, environmental effects, and mission conditions can change from moment to moment. The first challenge occurs when warfighters overstate their individual needs of a shared resource, leading to inefficient allocation. Overstatement may bring local optimization; however, it can cause global inefficiencies that result in decreased overall mission success. This challenge is addressed by using computational mechanism design, more specifically, the dynamic Vickrey-Clark-Groves allocation mechanism.

The second challenge involves resource availability that may change frequently. Such variability occurs in a wireless mesh network, where routes and bandwidth may vary over even small intervals of time. To address this challenge, an adaptive quality-of-service (AQoS) approach is taken, and the available resource is allocated using the Dynamic QoS-based Resource Allocation Model (D-Q-RAM). Computational mechanism design is used to allocate sensors, and the AQoS approach allocates the available network bandwidth in a way consistent with the sensor allocation, providing an approach for dealing with resource allocation and adaptation in a dynamic environment. This report presents the experimental results of applying this approach.

1 Introduction

When warfighting missions are conducted in a dynamic environment, the allocation of resources needed for mission operation can change from moment to moment. This report addresses two challenges of resource allocation in dynamic environments:

1. **Overstatement of resource needs:** A user may request allocation of a resource, motivated by self-interest for local gain that is an exaggerated expression of actual need. As a result, others may be denied the use of a resource, thereby adversely affecting overall mission success.
2. **Unpredictable network availability:** When the availability of a network is unpredictable, applications' capability for accomplishing missions is limited. The management of network bandwidth that must be shared among multiple applications is of significant concern in this work.

Both of the above situations represent examples of operations in a dynamic environment. One contributor to the environment's dynamic nature involves operational requests from the soldiers who need resources to complete a mission. A second contributing factor involves use of a wireless network, where it is not possible to predict the amount of network bandwidth available with sufficient certainty. We must consider both of these factors. When dynamic effects can affect operational mission utility, the ability to adapt resource allocation is a significant need.

We address these two challenges with a twofold approach. Resolving overstated resource allocation requests is based on computational mechanism design [Dash 2003]. To deal with unpredictable resource availability, notably bandwidth in a wireless environment, an adaptive quality-of-service (QoS) approach is taken [Hansen 2010]. Using these approaches together, we present a solution for dealing with resource allocation in a dynamic environment.

We examine the problem by using a scenario that is representative of resource allocation in current Department of Defense (DoD) deployments. We restrict ourselves to a case where the scope of resource sharing is limited to small groups of deployed soldiers. Initial experimental results suggest that the combination of computational mechanism design and adaptive QoS (AQoS) provides an effective approach for dealing with resource management.

This report is organized as follows: Section 2 provides background information regarding the scenario description as well as the approaches we use: namely, computational mechanism design and AQoS. Section 3 includes details of the scenario as well as the implementation of the approach. Section 4 presents experimental results for the scenario. A brief summary of the report appears in Section 5. Appendix A contains a description of the details of the scenario, and Appendix B contains a discussion of sample data used in the experiment.

2 Background

2.1 Scenario Approach

In this report, we will use a scenario, based on an operational configuration, to guide the approach and presentation of our results. First, we will describe an example that is representative of current operations and then we will introduce our experimental scenario that manifests similar characteristics. We will examine our experimental scenario in the remainder of this report.

2.1.1 Operational Configuration

Figure 1 illustrates the operational configuration that we considered. An aerostat carries a payload that typically consists of sensors. The purpose of using an aerostat is to achieve improved situational awareness through Intelligence, Surveillance and Reconnaissance (ISR). The aerostat is controlled through a radio link to the Ground Control Station (GCS). Various consumers can make requests, using a different wireless link, for information from the sensors, and the GCS manages those requests, which involves arbitrating where the sensors will be pointed.¹

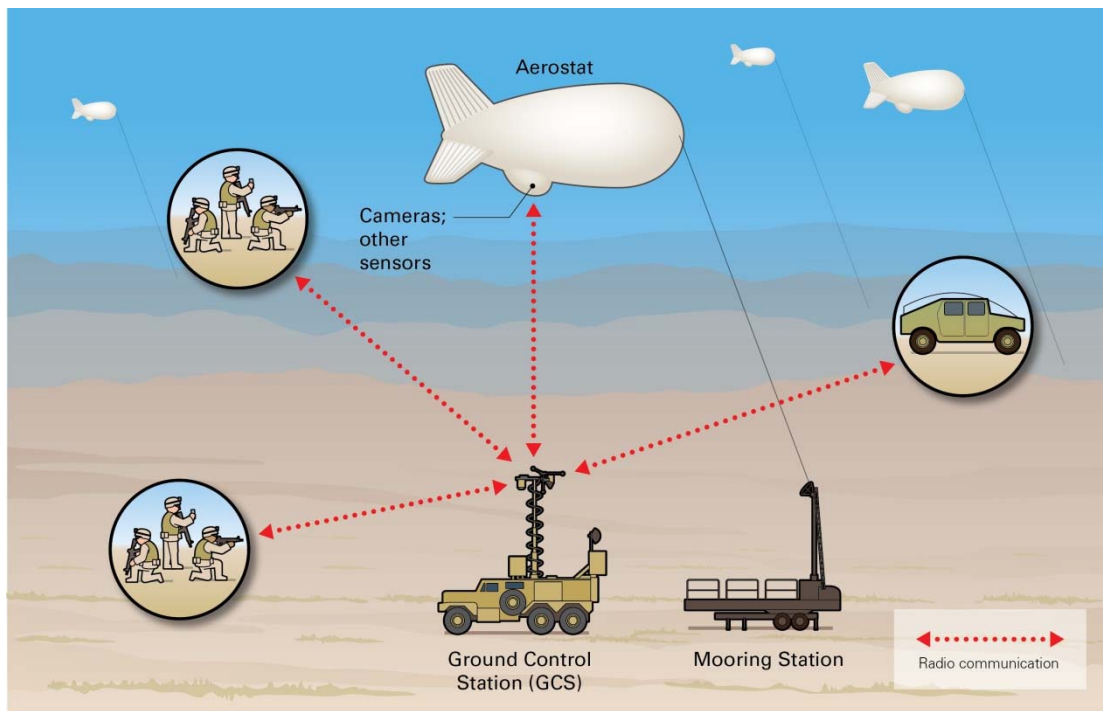


Figure 1: Operational Configuration

¹ Several sources illustrate the use of aerostats in current operations. These include discussions of a general nature (<http://www.wired.com/dangerroom/2011/04/cameras-spy-balloons-surge-in-afghanistan/>), the use of aerostats for distribution of sensor data (<http://www.wired.com/dangerroom/2011/01/all-seeing-blimp/>), the role of aerostats in tactical surveillance (<http://www.ground-combat-technology.com/gct-home/244-gct-2010-volume-1-issue-1-june/2939-tactical-surveillance.html>), dealing with sensor overload (<http://defensesystems.com/articles/2010/04/06/cover-story-sensor-overload.aspx>), and the use of aerostats for defense against cruise missiles (<http://www.globalsecurity.org/space/systems/jlens.htm>).

The operational configuration depicted here can have several variants that illustrate resource allocation concerns, including the following:

- **Physical characteristics:** An aerostat can vary in size and in the altitude at which it operates, both of which have implications for the employed sensors. In addition, the aerostat can be tethered from a tower or a mobile platform. Whether an aerostat is tethered or mobile can affect the nature of communication from the aerostat (e.g., whether it is a wireless or wired connection).
- **Payload:** The payload can include various sensors such as video cameras, infrared cameras, electro-optical sensors, or acoustic sensors. Each kind of sensor has implications for the amount of bandwidth necessary to transfer data as well as for other quality characteristics, such as image resolution. Furthermore, if a sensor is steerable, it has implications for the ability to share the sensor among multiple users; however, this will not be the case if a sensor is omnidirectional.
- **Coordination:** There can be multiple consumers of the available information, and it is necessary to coordinate the fulfillment of their requests. In the presence of multiple consumers, there will be issues about the order of delivery of information to those consumers.

A general concern in the use of aerostats is their scope of application. For example, an aerostat can be used to provide situational awareness to a Forward Operations Base (FOB). However, it is also possible to integrate information from multiple sensors on multiple aerostats to obtain an integrated view of an area. Thus, while there are several applications of an aerostat in operational scenarios, as Figure 1 illustrates, the experimental scenario we have chosen represents a typical application of an aerostat for ISR.

2.1.2 Experimental Scenario

In this section, we introduce the experimental scenario that we consider in the remainder of this report (see Figure 2). The scenario missions are to (1) defend the area against hostile attack, and (2) identify and capture potential hostiles that may approach the area. Effective allocation of resources, such as sensors and bandwidth, is necessary to increase the probability of mission success. In this report, we concentrate on the second scenario mission to illustrate the problem of resource allocation.




Figure 2: Overview of Experimental Scenario

The experimental scenario includes two squads deployed in an area, as Figure 2 illustrates. The squads (S_L and S_R) are on the left and right sides of the Tactical Operations Center (TOC), respectively. The TOC camera is a shared resource that may be allocated exclusively to either squad until the next allocation decision is made. In addition, the TOC camera will require bandwidth from a shared network to distribute video images.

The TOC has the ability to communicate with both of the squads. It also has the following capabilities:

- A Resource Management Agent (RMA) is responsible for allocating shared resources to the squads. In our experiment, the TOC camera is the shared resource that either squad can use. The RMA employs a computational mechanism design approach to allocating the TOC camera.
- A Facial-Recognition Server (FRS) can receive images from smartphones the squads use, compare them to a local database, and then return the probability that an individual matches a stored image

achieve the mission to identify and capture hostiles that may approach the base, the commander can use checkpoints to identify potential hostiles. If checkpoints are placed on the road, they are readily visible, and it is unlikely that the hostiles will approach the base. This eliminates opportunities to identify and capture hostiles. Instead of working from visible checkpoints on the road, each squad will be concealed behind existing structures or trees next to the road, so as not to deter the hostiles from approaching. Each squad will use the TOC camera to determine when cars are approaching and stop them for identification checks just past the structure or trees blocking the squad from view.

Each squad also has a Ground Vehicle Camera (GVC) that it can use to monitor traffic approaching the checkpoint. However, one disadvantage of using the GVC is that it requires placing the vehicle visibly on the road, and consequently may deter hostiles from approaching. Thus, using the GVC reduces the chance of capturing the hostiles. Each squad also has a Squad Camera that allows the commander in the TOC to obtain situational awareness of the operating area of a squad.

Use of the TOC camera allows the squads to monitor traffic approaching the checkpoint without the deterrence factor. However, the TOC camera can be allocated to only one squad. The squad that does not get the TOC camera will use the GVC as fallback. In addition, each squad has the ability to capture images with a smartphone and send them to a Facial-Recognition Server in the TOC to identify a potential hostile.

Appendix A presents the details of the scenario in a step-by-step format.

The experimental scenario has characteristics similar to the operational scenario that Figure 1 illustrates. In particular, there is a steerable high-resolution camera mounted on a tower; this is a shared resource among multiple squads and represents the sensor payload carried on an aerostat. The underlying network serves to communicate the information provided by the sensors to multiple squads, and the network characteristics (notably bandwidth) must be managed, especially when there can be dynamic effects on the network such as bandwidth availability or data loss due to environmental factors.

The similarity of the experimental scenario to an operational situation allows us to investigate the following questions:

Is it possible to provide an efficient means for resource allocation that prevents possible inefficiencies due to overstatement of a resource need?²

Is it possible to provide dynamic, satisfactory adaptation of resource allocation in situations where there is a degradation of resource availability?

In the following sections, we examine the above questions.

2.2 Resource Allocation in the Presence of Overstatement of Resource Needs

In the scenario, the commander wants to make allocation decisions that achieve the most effective use of the video sensors and network resources to accomplish the mission. In tactical operations such as the one described above, uncertainty, dynamism, decentralization, and interaction between the human and computational elements (such as resource allocation software, video and Facial-Recognition Servers, and quality of service management agents) of the system pose special challenges to allocating resources. Uncertainty comes from different sources: for example, how a hostile will act or how inclement weather might affect image capture, and whether a hostile will be identified. The situation is dynamic because either the mission can change or the warfighters must react to unplanned situations. Decisions require input from the warfighters in the field; that is, the input is decentralized and must be aggregated. In addition, since human input influences the allo-

² Although in this report we focus on the sensor allocation for this question, we plan to extend the techniques to discourage overstatements in network requirements as well.

cation of sensors and the network resources needed to operate them, there is interaction between human and computational elements of the system.

We can describe the problem of resource allocation in the tactical environment as an optimization problem, as follows: “How can resources be allocated in a way that maximizes expected mission value or success?” Even if we assume away the complexity of this optimization problem, one fundamental problem still remains. The commander must rely on decentralized reports from warfighters at the edge to acquire situational awareness and elicit resource needs; these reports, together with resource availability information, constitute the inputs to this optimization problem.

Since those who report their situations and needs are also directly affected by the resulting allocation decisions, the opportunity for misreporting arises. That is, warfighters may overstate their need for a resource or how useful a resource would be in their given situations. They may be trying to do what they think is best for their particular mission, but may unknowingly use resources that would be more beneficial for other warfighters. For example, warfighters have been known to overstate the precedence³ of messages with the intention of speeding up their delivery, but in doing so, they unintentionally saturate the network, reducing its effectiveness not only for them but also for others [US Army 1991]. In Iraq, units requested recurring aerial surveillance of the same nearby targets day after day, just to make sure they had an unmanned aerial vehicle (UAV) available in case they needed it [Greenberg 2010]. If the input to the optimization problem is incorrect because it is intentionally distorted (even when that is done with the best *local* intentions), the resulting resource allocation will not be optimal, and consequently, the commander will not get the most value out of the scarce resources.

2.2.1 Computational Mechanism Design

The problem of resource allocation among strategic agents⁴ (i.e., agents that will deviate from the desired social behavior if doing so is in their best interest) has been a topic of study in economics. For example, the Vickrey auction solves the problem of allocating a good to the agent that has the most value for it [Vickrey 1961]. In this auction, the auctioneer elicits the value each agent has for the good, which is private information that only each agent knows and could misreport. The good is allocated to the agent with the highest value, and that agent pays the second highest price. A key property of this auction is known as *incentive compatibility*, meaning that the best strategy for each agent is to report truthfully the value for the good. To understand this concept, consider an auction with three agents, *A*, *B*, and *C*, whose true values for the auctioned item are 8, 9, and 10, respectively. If agent *B* overstates its value as being 11, it does get the good, but must pay 10, achieving a payoff (true value minus payment) of -1 (i.e., a loss). This outcome shows how an agent can be hurt by reporting a higher value than is truthful.

Similarly, an agent cannot benefit from reporting a lower value than is truthful. For example, if agent *B* reports a lower value, this will not affect the payment if it wins, but doing so can change the outcome and result in the good being allocated to agent *A*. Consequently, the best strategy is to report truthfully. A mechanism such as the Vickrey auction defines a decision protocol so that some goal (e.g., maximizing the aggregate value agents get) is achieved even when agents pursue

³ *Precedence* is the term used in the military to indicate the urgency of a message.

⁴ The term *agent* refers to a participant in an auction and other economic mechanisms who makes decisions about how to act in that setting.

their own best interests (behave strategically). Computational mechanism design (CMD) is a discipline that studies the computational aspects of those decision protocols, and how game theory and economic principles of mechanism design are applied to multi-agent systems [Dash 2003].

The Vickrey auction, which was later generalized by Clarke and by Groves and thus is known as the Vickrey-Clarke-Groves (VCG) mechanism, handles a single allocation decision [Clarke 1971; Groves 1973]. However, in tactical environments, it is unlikely that a single allocation decision will remain optimal over time because the needs of the mission change continuously. If resource allocation cannot keep pace with the dynamism of the tactical environment, the resulting suboptimal resource allocation will affect mission success and could also affect the behavior of the warfighter. If warfighters think that the resource-allocation approach cannot respond to changing situations promptly, they may be inclined to request resources just in case they need them when the situation changes [Greenberg 2010].

We use dynamic VCG to address possible resource overstatement by a user in the presence of dynamism and uncertainty [Bergemann 2006; Cavallo 2009]. In dynamic VCG, agents report not only the value they have for different outcomes, but also how different decision outcomes change the probability of transitioning to different states. The formalism used to represent this kind of report is a Markov Decision Process (MDP). Figure 3 shows a sample MDP for allocation of a video stream. The possible states that the agents could be in are represented by the circles. The thick arrows are the possible allocation decisions that could be made, and the arcs represent the probability of transitioning to a new state, given that the agent is in a particular state and a particular allocation decision is made. An MDP also encodes the value that the agent earns when each state is reached (which is not shown in Figure 3). Obviously, it is not practical that a warfighter report an MDP. We address this issue in Section 3.2.1, but for the rest of this section, we assume that agents report MDPs.

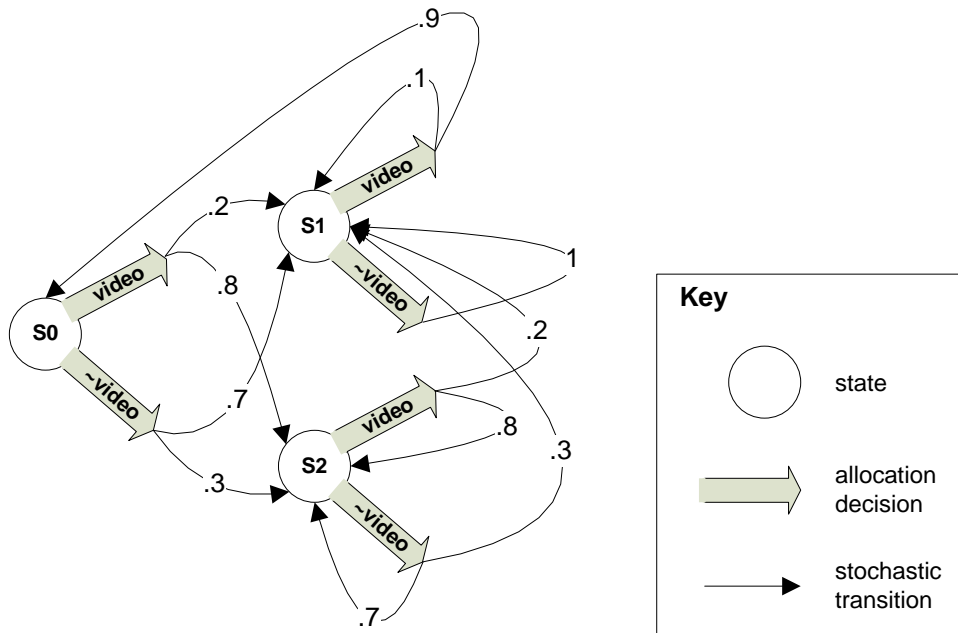


Figure 3: Markov Decision Process

Unlike in other resource allocation approaches, in dynamic VCG agents report their MDPs instead of reporting a direct request for resources. With this approach, agents are effectively reporting how different allocation decisions would affect their chances of reaching different states. An optimal resource allocation can be computed by creating a joint MDP based on the individual MDPs and determining the decision that maximizes expected mission value over time, given the current state of the mission. The joint MDP allows us to handle not only the uncertainty of the tactical environment (by including the stochastic transitions), but also its dynamism; it evaluates how allocation decisions affect not only the current state but also subsequent expected states of the mission.

Since dynamic VCG is incentive compatible in expectation, an agent maximizes the expected net value it gets by truthfully reporting the MDP. Also, allocation decisions are not static, but repeated over time; thus, agents do not need to overstate needs in anticipation of future resource allocations they may require. Instead, if they enter a state where, if having a resource would provide the most benefit to the mission over time, they would receive that resource.

Dynamic VCG achieves incentive compatibility by having agents make payments computed in a way such that each agent maximizes its expected payoff (i.e., value obtained from assets and resources minus payments) by truthfully reporting the MDP. Although in traditional markets these payments are made in some currency, other forms of payments are possible, as we will discuss in Section 3.2.1. The payment is computed as the expected value an agent takes from the rest of the agents by participating in the mechanism. For example, if a sensor is allocated to agent *A* based on the reports by all the agents, that sensor cannot be used by other agents during the period of time *A* has it. Suppose agent *B* could achieve some value with that sensor. The first observation is that the expected value agent *A* can achieve with the sensor is higher than agent *B*'s; otherwise, the efficient allocation would have allocated the sensor to agent *B*. The second observation is that if agent *A* makes a payment equal to the expected value agent *B* could have achieved with the sensor, agent *A* will have a positive payoff in expectation. In dynamic VCG, the best strategy for an agent is to report truthfully [Cavallo 2009].

2.3 Adaptive Quality of Service (AQoS)

AQoS is a technique that enables applications to change their own service levels in response to available resources [Hansen 2010]. Our approach to AQoS is a technique called Distributed QoS Resource Allocation Model (D-Q-RAM). D-Q-RAM employs a distributed optimization heuristic that enables application adaptation in such a way that resources are near-optimally⁵ allocated. Optimality is judged in terms of utility values that are assigned to different levels of service for each application. An allocation is considered optimal if it is not possible to reallocate resources in such a way that the sum of the utility values across all applications can be increased.

2.3.1 QoS Resource Allocation Model (Q-RAM)

D-Q-RAM is a distributed version of an earlier technique called Q-RAM [Hoover 2001]. Q-RAM is a QoS optimization heuristic that has been successfully applied to a wide range of applications over the past decade. Applications have ranged from multimedia applications to phased array ra-

⁵ We use the term "near optimal" because of the discrete nature of the problem. That is, it may be that optimality could be satisfied in a continuous case, but is only near optimal for a discrete representation.

dars. An application's QoS needs are modeled as a set of *setpoints* in a multi-dimensional QoS space. Each QoS dimension represents one aspect of application quality. For example, frame rate, image resolution, and compression level might be three QoS dimensions for a video stream. For each setpoint, there is an associated utility value describing a resource requirement and its usefulness to the user.

The basic problem solved by Q-RAM is that of choosing from a set of setpoints Q_i , an *active setpoint* $q_i \in Q_i$ for each application i in a set of applications, so that the sum of the application utilities $u_i(q_i)$ is maximized while the sum of the application resource requirements $r_i(q_i)$ does not exceed a specified resource limit. In general, this is an NP-hard problem, but Q-RAM applies heuristics to find an approximate solution in polynomial time.

The key idea behind the Q-RAM heuristics is to allocate resources while satisfying the Kuhn-Tucker condition [Kuhn 1951]. The Kuhn-Tucker condition applies to the continuous case, and we will later generalize this to the case with discrete setpoints. Given a set of N applications with continuous concave utility functions $u_i(r_i)$ where r_i is the amount of resource allocated to application i , and a maximum resource capacity R_{max} , Kuhn-Tucker tells us the optimal allocation will occur when

1. All of the resources have been allocated among the applications. That is, $\sum_{i=1}^N r_i = R_{max}$.
2. The slopes of the utility curves (known as the *marginal utility*) of each application are equal at the allocation setpoints. That is, $u'_1(r_1) = u'_2(r_2) = \dots = u'_N(r_N)$.

It is also possible to include a value of weight for a flow where a flow has a source and destination address, defined as a pair (IP address, port identifier). The value of weight indicates the relative importance of a flow and may be between zero and one. The larger the value of weight, the more important a flow is relative to other flows. Note that if a weight has a value of zero, that does not mean the flow will not get any resources; rather it indicates a flow will operate at its lowest QoS level.

2.3.2 Centralized Q-RAM

Past implementations of Q-RAM have used a centralized approach to satisfying the Kuhn-Tucker conditions [Hoover 2001]. In this approach, there is a centralized resource manager that maintains profiles for all application types. Each application profile consists of the set of setpoints comprising a resource requirement and a utility value. Optimization is achieved by initializing the setpoints for all applications to the lowest setpoint, then repeatedly choosing the next application setpoint for which the slope (increase in utility versus increase in resource) is highest. If the utility and resource values of application setpoints are reasonably well distributed, this heuristic will result in an allocation where the slopes at the chosen setpoints are nearly equal, making it an approximation to the Kuhn-Tucker condition.

The weaknesses of centralized Q-RAM for application to tactical settings that involve local scope and a dynamic environment include the following:

- The capacity of each resource must be known and constant in time.

- Q-RAM must have access to the utility curves for all applications.

2.3.3 Single Hop D-Q-RAM

The key idea of D-Q-RAM is to satisfy the two optimality conditions described above in a distributed rather than centralized manner [Hansen 2012]. In this report, we describe only the case of a single hop. We will use the example shown in Figure 4 to illustrate our approach. In this example, sources S1-S3 are video sources sharing a wireless output link on a router R . The utility of a flow is described by using a function of only the bandwidth consumed by the flow. It is assumed that the links from each source to R are wired and are of sufficient capacity.




Figure 4: D-Q-RAM with Single Router

In D-Q-RAM, the utility curves are not centrally managed, but are private data managed by each application. This results in much better scalability than is offered by a centralized approach. In our example, S1 receives the most utility for a given amount of resource, while flows S2 and S3 have utility curves that are weighted with weight factors between zero and one. Because of the weights we have chosen in our example, when each application operates at a setpoint with equal slope, S1 receives the largest resource allocation and thus operates at a higher QoS level than the other flows.

To satisfy Kuhn-Tucker Condition 1, we must ensure that the router is consuming all, or nearly all, of its resources. While it is difficult to know the exact capacity of a wireless link, it is possible to know when the link is oversubscribed by monitoring the output queue on the router R . If the link is oversubscribed, the output queue will grow large. In our approach, when the queue length stays above a threshold L_{TH} over a period of time T_O , a control message (described in the following paragraph) is sent to all sources using that router, instructing them to operate at a lower QoS level and thus reduce the demand on that link. If the queue length has remained below the threshold for a time period T_U , then a control message is sent to all sources instructing them to operate at a higher QoS level, thus increasing the demand on the link. These two control messages

establish a feedback control loop with demand from the applications being increased and decreased so as to maintain the link just below the point where it becomes oversubscribed.

To satisfy Kuhn-Tucker Condition 2, we need to ensure that the marginal utility at the operating setpoint⁶ for all applications is equal or near equal. We approximate this by having the router maintain a current marginal utility value m that is periodically transmitted back to the sources⁷ (currently once per second). When the oversubscription condition holds, R increases its *marginal utility* value and transmits it back to the sending applications S1, S2, and S3. This causes them to operate at yet a lower resource-demand point on their utility curves. When the undersubscription condition holds, the m value is decreased, allowing applications to operate at a higher resource-demand point on their utility curves. In our current implementation, we use a message to increase or decrease the value of m . The result is that an application flow is changed over time.

Figure 5 provides a description of the algorithm for a router distributing Marginal Utility messages.

⁶ An *operating setpoint* is a setpoint that lies on the concave majorant, as illustrated by the curves in Figure 4, and is being used by an application at a particular time [Lee 1999].

⁷ A *Marginal Utility* message is transmitted periodically because we chose User Datagram Protocol (UDP) as the transport protocol, which does not provide reliable delivery. We chose UDP, as opposed to Transmission Control Protocol (TCP), because the unpredictability of a mesh network would cause difficulty in maintaining a TCP connection.

```

do every 250 ms
  call Update Marginal Utility Value
  if marginal utility has changed or
    last Marginal Utility message sent > 2s ago then
    call Send Marginal Utility Messages
  end if
end do

proc Update Marginal Utility Value
  if no queue length change for 1s then
    marginal utility := 0.0
  else
    if queue length > high mark then
      increase router's marginal utility value
    end if
    if queue length < low mark then
      decrease router's marginal utility value
    end if
  end if
end proc

proc Send Marginal Utility Messages
  for each active flow through this router do
    create Marginal Utility message containing m
    send Marginal Utility message to flow source
  end for
end proc

```

Figure 5: Algorithm for Distributing Marginal Utility Messages on a Router

As indicated in Figure 5, there are a number of parameters used in the algorithm. These include

- the execution rate of the main algorithm (250 ms)
- the minimum rate at which a Marginal Utility message may be sent (2 s)
- the quiescent period for resetting the marginal utility value (1 s)
- the upper (lower) queue length for increasing (decreasing) the value of marginal utility
- the expiration time for which a flow is active (2 s)

On receiving a marginal utility value from the router, each source chooses to operate at the set-point from its privately held utility curve such that the slope is closest to but not exceeding the received value. This ensures that all sources are operating at or near the same marginal utility. In combination with the feedback-control mechanism ensuring that resource consumption is near the

maximum, both Kuhn-Tucker conditions are satisfied (or approximated) and thus the resulting allocation is near optimal (to within one setpoint). Note that the procedure *Update Marginal Utility* initially checks for a change in the queue length for the last second. This is a check for quiescence from the source.

2.4 Integrating the Approaches

Having described the two aspects of resource allocation, namely, one for the TOC camera and the other for network bandwidth, we proceed to address their integration. The RMA allocates the TOC camera using CMD by communication with warfighters through the exchange of messages, as described below. Bandwidth is also managed by D-Q-RAM to support image distribution from the various cameras in the scenario.

The integration of overall resource allocation is based on the following:

1. A user (such as a soldier in a squad) requests allocation of the TOC camera by specifying the information relevant to resource allocation (Section 3.2.1), such as the probability of a hostile approaching.
2. With that information, the RMA determines the user to be granted control of the TOC camera.
3. When the RMA allocates the TOC camera, it assigns a weight value for the utility of a TOC camera video flow, and the weight values are used for managing network traffic by AQoS. A higher weight indicates that the user will be granted a higher proportion of the available network.
4. The RMA distributes new values of the weights to those application components that are the source of a flow. These flows are managed by AQoS.
5. Successive message traffic is based on the weighted utility curve that applies to the resource allocation.
6. As the mission progresses, if the available network bandwidth is insufficient to meet the demand, AQoS notifies applications that the network is oversubscribed. The applications will then select another setpoint that will decrease its bandwidth use and dictate its future network transmission characteristics.

Thus, integration of resource allocation involves two facets: allocation of the TOC camera and allocation of the network bandwidth needed for the distribution of TOC video images. The allocation of the resource itself is achieved by the use of the CMD approach. Given that a resource has been allocated, bandwidth allocation is achieved by the use of AQoS. We believe this represents a consistent approach to allocating available bandwidth and sensors. However, further examination of an integrated approach is required. For example, there are open questions regarding whether our integrated approach is incentive compatible or near optimal.

3 Scenario Implementation

3.1 Configuration

The experimental scenario was conducted at an airfield at Camp Roberts, a California Army National Guard base. The TOC was located in a trailer on the airfield, Squad Left (S_L) was on a dirt path near the airfield, and Squad Right (S_R) was on the road leading away from the airfield. An aerial view of the area showing the deployment of the components of the experimental scenario is shown in Figure 6. This figure represents a realization of the earlier description of the experimental scenario for which Figure 2 provides a preliminary depiction.




Figure 6: Aerial View of Scenario Configuration

The squads, as well as the TOC, are all connected by a wireless mesh network⁸, and each site contains a mesh radio. Connectivity between wireless radios involves only a single hop to the TOC. The radio for the TOC is mounted on an antenna mast next to the TOC and connected by a wired network to the steerable camera mounted on the mast. To focus our study on resource management, we do not address communication between the squads (via the TOC).

Figure 7 presents a diagram of the network components, denoting the two squads as well as the TOC. Each squad has a Squad Camera, a Ground Vehicle Camera (GVC), and a laptop that uses a wired connection to communicate with a wireless mesh radio. In addition, each squad has a smartphone that also communicates over the wireless network. In the TOC there are two laptops:

⁸ A *mesh network* is a network where, in addition to receiving and sending application packets, each node serves as a router.

one is used for the TOC video display, TOC video server, and FRS server and the other for communication with the RMA. The wired network in the TOC and the squads has a maximum bandwidth of 100 Mbps.

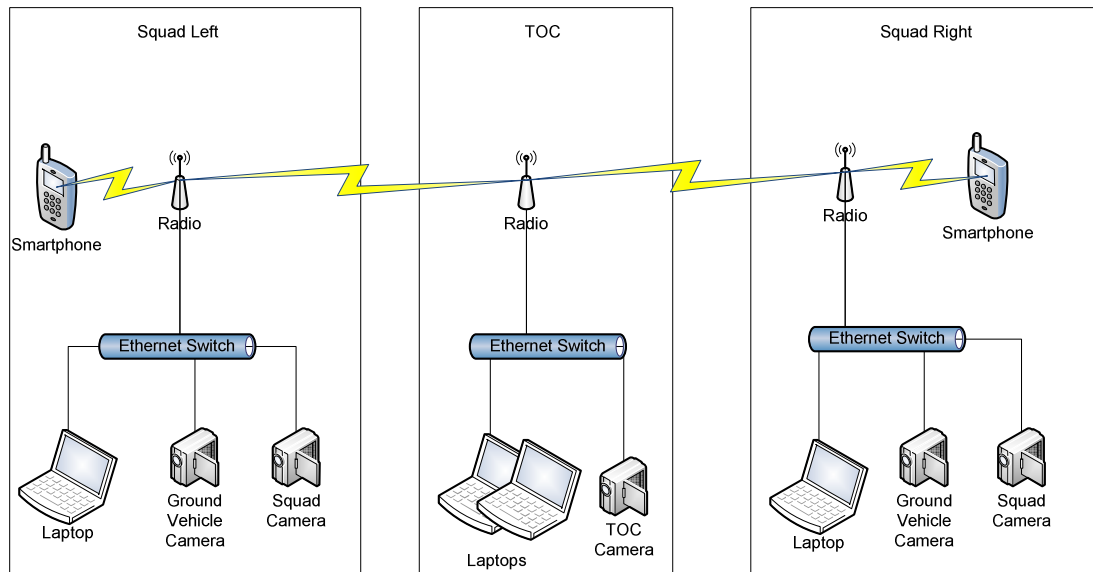


Figure 7: Logical Layout of Components of Wired and Wi-Fi Network

3.1.1 Tactical Operations Center Configuration

In the center panel of Figure 7, we see the TOC, which is responsible for the overall conduct of the mission to detect and capture hostiles who may approach the base. The TOC controls resource management of the TOC camera and distribution of TOC camera images to the squads, as well as managing the facial-recognition processing. The RMA determines the resource management of the TOC camera, a shared resource, by using Dynamic VCG that we described in Section 2.1.1. Video images of variable quality from the TOC camera are transmitted to the squads, with a quality-of-service level based on the weighted utility curves that we described in Section 2.3. The purpose of the Facial-Recognition Server (FRS) is to identify an individual from an image it receives from a squad smartphone.

The TOC can provide a display from multiple cameras as well as information regarding facial recognition. Figure 8 illustrates a typical display. Two adjacent images represent the images associated with each squad, each a two-by-two grid representing the squad on the left and on the right, respectively. The images at the top represent images from the squad camera and GVC. However, where one of the squads was using the TOC camera, the image from the TOC camera displays in place of that squad's GVC video.

The images at the bottom of each two-by-two grid are related to facial recognition. The image on the left is the image that a squad smartphone sends to the FRS, while the adjacent image shows the submitted image with a bounding box denoting the area of the image that will be used in processing the image. If the FRS finds a match, it returns the identity of the individual and displays it at the bottom of the image. If the FRS finds no match, it will return an indication of this result and

no bounding box will appear over the image. In the event that the FRS exceeds a timeout for its response, a retry is attempted.



Figure 8: Screenshot from TOC Display

In addition to the display, the RMA collects input from the squads regarding resource allocation requests and to use in determining the allocation.

3.1.2 Squad Configurations

Each squad, shown in the side panels of Figure 7, employs a mesh radio that wirelessly connects to a smartphone and to the TOC radio. The radio was also connected over a wired network to a Ground Vehicle Camera, a squad camera, and a laptop. This was possible because these items were co-resident in a vehicle.

The laptop in each squad runs an application that allows the squad to report the information needed to compute the allocation of the TOC camera and to receive notification of the result of the allocation request. Figure 9 provides a screenshot of a typical warfighter request for resource allocation. We will describe the details for the application in Section 3.2.1. The resource allocation application on the squad laptop communicates with the RMA at the TOC over the mesh network.⁹

⁹ Normally, a squad would report its resource requests to the RMA. However, for logistics reasons, in the experiment the resource requests were entered by an individual who was in the TOC.

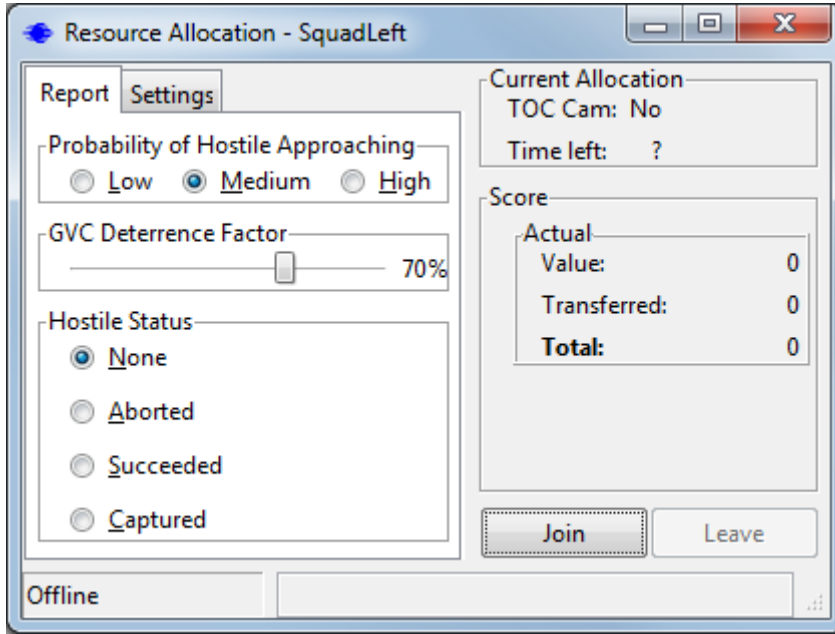


Figure 9: Squad Interface for Requesting Resource Allocation

Facial recognition is accomplished by the use of smartphones. The smartphone can capture an image of a subject and, if needed, modify the quality of the image to be suitable for the available bandwidth by using a weighted utility curve. The images are then sent to the FRS in the TOC for processing.

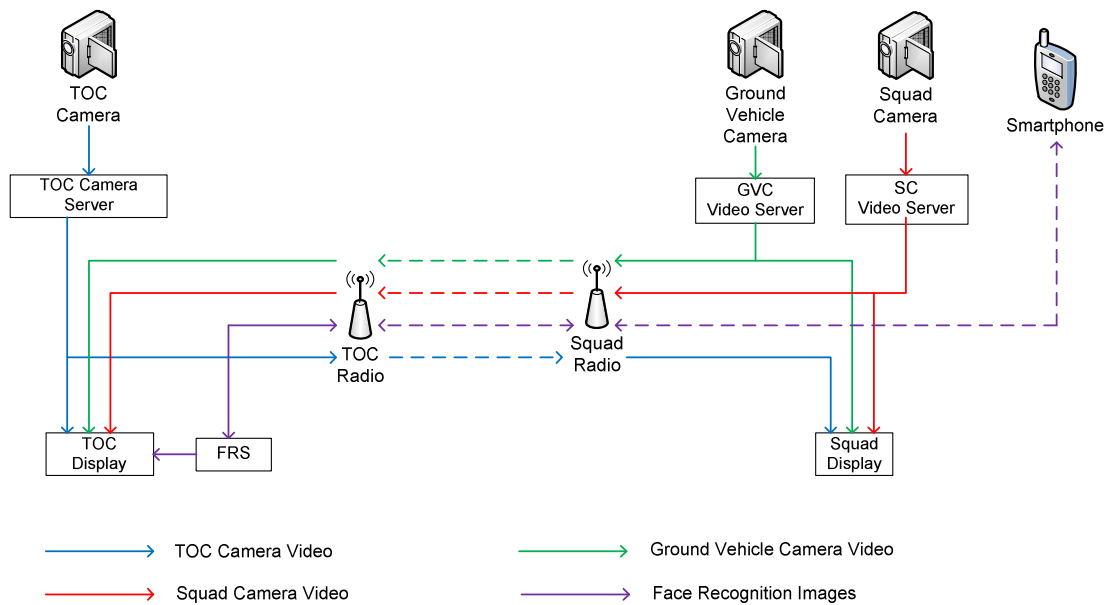
3.1.3 Information Flows

In the experiment are several network flows, introduced in Section 2.1.2. Each flow has a utility curve as well as weights, indicating the importance of a flow. The flows use messages that are grouped into the following categories:

- control messages: messages used to control the adaptation process (e.g., *Marginal Utility messages* and *AQoS Weight messages*)
- AQoS managed messages: GVC video, squad video, TOC camera video, and FRS images
- Resource Management messages
- others: for example, a background HTTP transfer

A control message conveys, either implicitly or explicitly, a behavior that a recipient is expected to perform. A *Marginal Utility message* is a message that contains the value of the marginal utility as the marginal utility was defined in Section 2.3.1.

Table1 summarizes the flows. In the table, there is no distinction between the squads on the left and right of the area, as their functions are symmetric (although only one squad at a time has access to the TOC camera). Also, some of the flows are bi-directional; for example, a squad can request allocation of the TOC camera, and a response will be returned. Such bi-directional flows are not described in the table.



Solid lines denote traffic over a wired connection and dashed lines denote traffic over a wireless link.

Figure 11: AQoS Managed Flows

The integration of CMD for camera allocation and bandwidth management was achieved by using messages. Overall arbitration of network resources is achieved through exchange of prioritized messages of various types. Table 2 represents the arbitration approach for messages that are transmitted from a wireless router where the various types of messages are prioritized according to their levels of importance.

Table 2: Relative Importance of Network Messages

Message Types	Importance
Control messages: <ul style="list-style-type: none"> • Marginal Utility messages • AQoS Weight messages 	High
AQoS controlled messages: <ul style="list-style-type: none"> • GVC video • squad video • TOC camera video • FRS images 	Medium
All other messages, including Resource Management Messages ¹⁰	Low

Thus, for example, a control message will be sent at higher priority than any AQoS controlled message, which will be sent at higher priority than any other messages. If there are multiple messages of a particular type awaiting transmission, they are queued in a first-in, first-out manner.

¹⁰ In general, when a network is heavily loaded, this can delay Resource Management messages due to the presence of higher priority messages. That case could not occur in this experiment, as Resource Management message functionality was entered by an individual in the TOC.

3.2 TOC Camera Allocation

The TOC camera was a shared resource available to each squad as part of this experiment. There are three basic steps required to determine and enact the most efficient allocation of the TOC camera.

1. Elicit relevant information from the squads.
2. Compute the allocation using the elicited information.
3. Enact the allocation.

We describe these steps in the following sections, and Figure 12 depicts the interactions among the different elements.

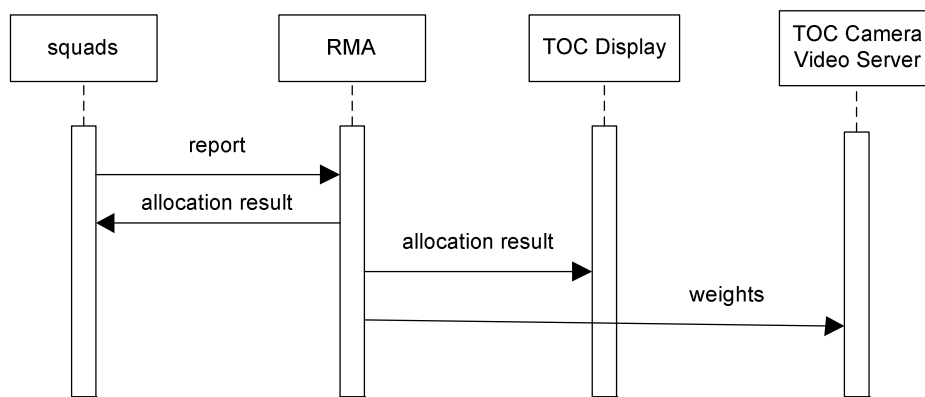


Figure 12: TOC Camera Allocation Sequence Diagram

3.2.1 Eliciting Relevant Information from Squads

In this experiment, the dynamic VCG mechanism that we discussed in Section 2.1.1 was used to determine the efficient allocation of the TOC camera. As we previously stated, dynamic VCG requires that each squad report its state and how different allocations—now and in the near future—could affect the probability of reaching other mission states. Although this information is encoded in an MDP, it would be impractical to have warfighters build an MDP to report mission status and needs. The approach we take is to have a template MDP suitable for the mission and ask warfighters for only the information necessary to fill in the blanks in the template. Warfighters can enter this information through a simple user interface, leaving the details of the MDP behind the scenes (that is, warfighters never see or have to know about the MDP) as part of mission planning (see Figure 9).

In our experiment, there is a single mission for each squad and consequently only a single template MDP was needed. In more realistic operations, this MDP template could be crafted during mission planning or it could be accessed from a database containing MDP templates for different types of missions. Figure 13 shows the initial MDP template for each squad. (Figure 14 will present a simplified version.)

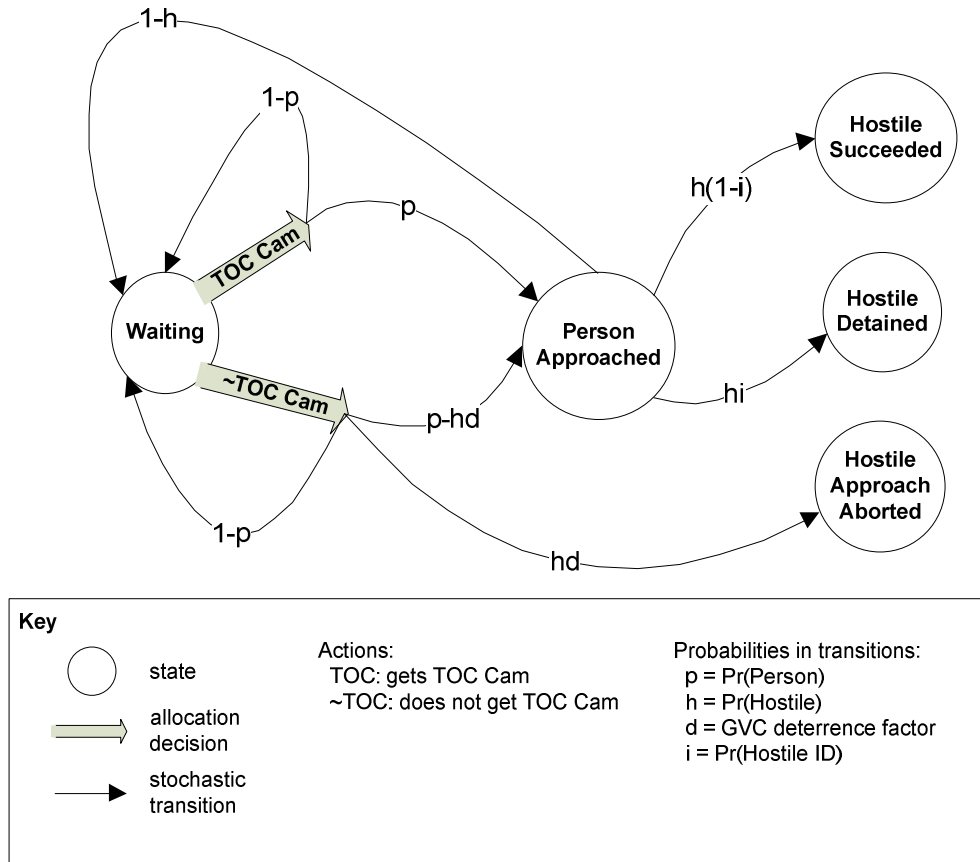


Figure 13: Initial MDP Template

Table 3 shows the parameter, its meaning, and the source that provides the value of the parameter.

Table 3: Parameters Used in the MDP

Parameter	Meaning	Source
mission end states	set of states that describe the end of the mission, including Hostile Succeeded, Hostile Detained, and Hostile Approach Aborted	mission planning
$\text{Pr}(\text{Person})$	probability that a person (hostile or not) will initiate an approach to the squad location	squad
$\text{Pr}(\text{Hostile})$	probability that a person approaching the squad location is a hostile	squad
$\text{Pr}(\text{HostileID})$	probability of positively identifying a hostile, when it is stopped at the checkpoint, by the use of facial recognition	system parameter
GVC deterrence factor	probability that the use of the GVC will deter a hostile from proceeding with the approach	squad

Notice, in Figure 9, the connection between the MDP and the parameters defined in the interface that the squad uses. The commander assigns the values of the end states of the mission. The allocation of assets will be computed so that it maximizes the expected mission value, following the commander's intent. Nevertheless, for the incentives to work, these values must represent some-

thing each squad would want to maximize individually. This could be, for example, time off after it completes the mission, or some score that is used to evaluate the squad in an after-action review. Table 4 shows the values used for the experiment.

Table 4: Mission End-State Values

End State	Value
hostile captured	5
hostile aborted approach	0
hostile succeeded with attack	-5

In our experimental scenario, we had one person approach each squad with certainty, so the probability of a person approaching, $\text{Pr}(\text{Person})$, was equal to 1. We also kept the probability of a hostile identification, $\text{Pr}(\text{HostileID})$, fixed at 0.80, representing the use of the smartphone for facial recognition.¹¹ In our scenario, we used some weight sets¹² that gave better QoS for the face-recognition information flows to the squad that was granted use of the TOC camera. In doing so, we introduced a dependency between resource allocation and the probability of positively identifying a hostile. Since the weights were not used in solving the MDP, this dependency was not taken into account for the optimization.

The GVC deterrence factor was part of the required input from the warfighter, because, depending on how open the area is surrounding the ground vehicle, this factor could change. Nevertheless, during our experiments, we kept this factor at its default value of 0.70. We elicited $\text{Pr}(\text{Hostile})$ as low, medium, or high, translating it internally to probability values 0.2, 0.5, and 0.8 respectively.

In the MDP shown in Figure 13, there is a state representing that a person has approached the squad. In general, squads would have to report every change of state to the resource allocation manager so that, if needed, resources could be reallocated based on the state change. In this MDP, however, no allocation decision is made in the state when a person has approached. In addition, no value is accrued by a squad for reaching this state. To avoid requiring squads to report every time a person approached them, the MDP was simplified by removing the state representing when a person has approached and computing the probabilities of reaching the other states directly from the waiting state. Figure 14 presents the simplified MDP template. The user interface for defining information for an MDP was as Figure 9 illustrates. Note that the GVC allocation did not involve the CMD approach. Instead, the assumption was that the squad that did not have access to the TOC camera would use its GVC as a fallback.

¹¹ Note that the value of $\text{Pr}(\text{HostileID})$ was independent of having access to the TOC camera. This is true at the sensor-allocation level, but may not necessarily hold when we take into account network bandwidth allocation. In other words, the value of $\text{Pr}(\text{HostileID})$ depended only on the ability of the FRS to identify a hostile and not on any property of the network, such as bandwidth.

¹² The various information flows can have a weight, indicating their relative importance. A *weight set* denotes a collection of weight values for a set of flows. These are discussed further in Section 3.3.

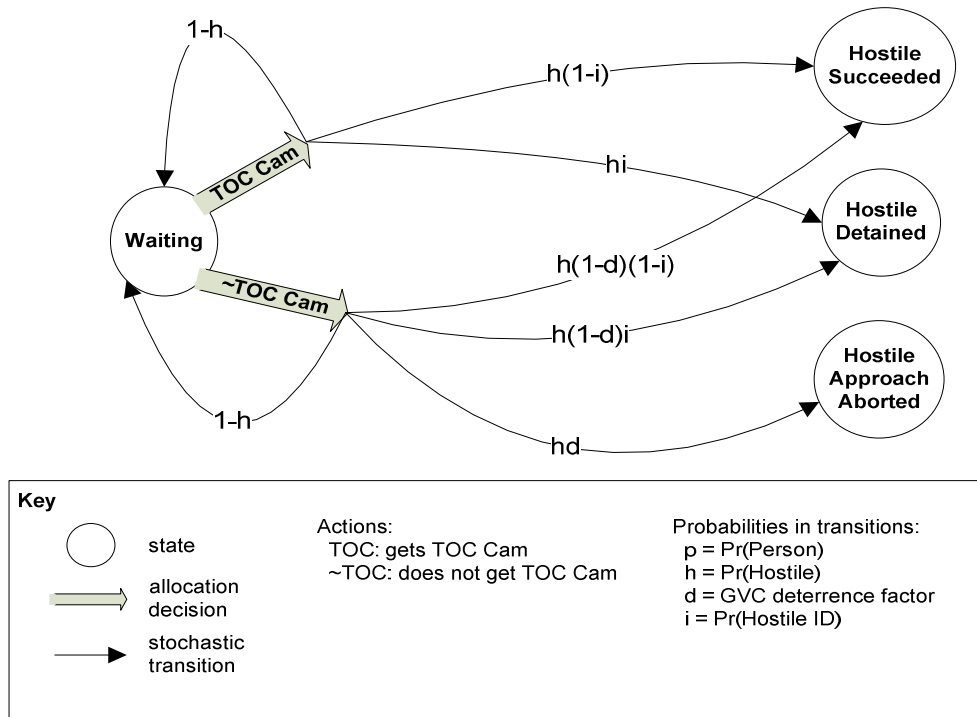


Figure 14: Simplified MDP Template

3.2.2 Computing Asset Allocation

The resource allocation is computed by the RMA executing in a computer at the TOC. The RMA can operate in two modes, automatic decision interval or manual. In the automatic mode, an interval is specified (e.g., 10 minutes) and an asset allocation is automatically conducted at every such interval. In the manual mode, an operator at the TOC decides when an allocation should be made. A third mode that we did not implement, but that could be used in operations, would involve computing a new asset allocation every time the squads change the input used to compute the allocation.

Every time an allocation is going to be computed, the RMA requests the squads' resource allocation application for the current input the users have specified. This information is sent to the RMA encoded as an MDP. The RMA then proceeds to compute the optimal TOC camera allocation. The computation entails creating a joint MDP with all the squads' MDPs, making the decision, and determining the optimal policy (i.e., which allocation should be made in each joint state to maximize expected value). The allocation decision corresponding to the current joint state in the optimal policy is chosen as the new asset allocation.

For experimental purposes, the allocation can be computed with or without use of the Dynamic VCG mechanism. If the mechanism is enabled, the RMA must compute the payments, if any, that each squad must make. These payments are in the same unit as the values determined by the commander for the different mission end states. If the units for those values are days off, for example, then a payment of -0.25 means that a quarter of a day is debited from the squad leader's

time-off balance. To compute the payments, the RMA must compute the allocation, removing one agent at a time, so as to determine the impact of the removed agent on the allocation outcome. On the other hand, if the mechanism is disabled, an agent can overstate his or her needs without suffering an adverse consequence. However, this overstatement will result in a decrease in the overall value for the commander (the social welfare). Section 4.2 provides experimental results for these two cases.

3.2.3 Enacting the Allocation

Once the allocation and the payments (if included) are computed, the allocation must be enacted. This involves carrying out several notifications. Each squad's resource allocation application receives a message with the allocations result, including whether or not the squad was granted access to the TOC camera and the payment (if any) the squad must make.¹³ The application shows this information on the squad's laptop. Another message is transmitted to the TOC display so that it can configure the view of the different video streams appropriately, as Figure 8 illustrates. Finally, messages are sent from the RMA to each video streaming application, indicating the weight that should be applied to the AQoS utility curve, corresponding to the different video flows.

3.2.4 AQoS as a Function of Weight

As we discussed in Section 3.1.3 and described in Table 1, there are multiple flows in the experiment. Some of those flows involve transmission over a wireless link. Due to the unpredictable nature of such a link, the routers manage those flows through an AQoS approach.

Depending on the outcome of the asset allocation, the importance of the different flows in the network will change. The relative importance of a flow is reflected through different weight sets that are selected based on the winner and loser of the TOC camera in the allocation. For each squad, there are four flows involving a wireless link that are assigned weight values: GVC, SC, Face, and TOC. The GVC and SC video flows originate at the squad and transmit to the TOC. For the TOC camera, the stream originates at the TOC and is transmitted to the squad that received the TOC camera allocation. For the Face flows, each squad sends images to be tested to the TOC FRS, and receives a response in return. While not all flows are used all the time in the scenario (e.g., the flow for the TOC camera to the losing squad is not always used), for the purpose of our experiment we keep all flows running all the time and use weights to limit the impact of the unused flows.¹⁴

One of the goals of the experiment was to investigate the application of weights to denote the relative importance of a flow. In particular, we are interested in the variation of QoS as a function of weight. For our experiments, we used four different weight sets for the flows shown in Table 5.

¹³ Only the accounting of payments is conducted during the mission.

¹⁴ Note also that the flow from a squad to the RMA for allocation of the TOC camera also involves transmission over a wireless link. This flow has not been included here due to logistics considerations. That is, rather than have an individual associated with a squad provide resource-allocation requests over the wireless link to the TOC, the resource requests were provided by an individual in the TOC for each squad. Additionally, the resource-management flows are not expected to consume a large fraction of the available bandwidth compared to, for example, a video stream.

Table 5: Weight Sets Used in Experiment

Weight Set	Outcome	GVC	SC	Face	TOC
1	Winner	1	1	1	1
	Loser	1	1	1	1
2	Winner	0.1	0.7	1	1
	Loser	0.8	0.7	0.7	0.1
3	Winner	0	0.7	1	1
	Loser	0.6	0.5	0.5	0
4	Winner	0	0.7	1	1
	Loser	0.3	0.2	0.2	0

For each weight set listed, Table 5 shows the weights applied to each of the flows for the winning and losing sides, where the winning side is the squad that receives allocation of video from the TOC camera. Weight Set 1 is included as a control in which all flows are equally weighted. The other weight sets include values that are specific to an experiment, discussed below.

The TOC camera is intended for assignment to the squad with the highest probability of having the hostile approach and be detected, thereby maximizing mission success. For this reason, in Weight Sets 2-4 the winner of the allocation decision is assigned a weight of 1.0 for the TOC camera, while the loser is assigned a weight of 0.1 or less. In our experiment, we expect the very low weight for the loser to allocate resources to that flow only if all other flows are maintaining high utility. We also expect that assigning a low weight to the loser will make those resources available to the winner, improving the winner's QoS. We will verify this by comparing the measured time-averaged utility (TAU) score, denoted τ , for the equal weight case (Weight Set 1), with the other weight sets. For a single video flow τ_j is defined as

$$\tau_j = w_j \frac{1}{T} \sum_i u_{j,i}$$

where w_j is the weight for flow j , T is the time period over which data on the flow was collected, and $u_{j,i}$ is the utility of the i -th frame received by flow j . A frame must be received to receive benefit, so frames that are dropped due to congestion in the network do not provide any utility in calculating τ .

The GVC is intended as an alternative to the squad that does not receive the TOC camera. For this reason, the loser is assigned a higher weight than the winner of the TOC camera. However, since the GVC is allocated on the side where we do not expect the insurgent to arrive, it is still assigned a lower weight than the TOC camera for the winner, with widening disparity for higher numbered weight sets. We expect these settings will provide reasonable but not high quality for the loser of the TOC camera and low quality for the winner who is not using this flow.

We assign weights for the squad camera flows to reflect that these flows are secondary to and require fewer resources than either the TOC camera or the GVC. In Weight Set 2, we weighted these flows just a little less than for the GVC and, for Weight Sets 3 and 4, we further reduced the weight on these flows to free more resources for the primary surveillance activities. As a result, we expect the SC flows to have QoS levels that are slightly lower than for the GVC, with left and right about equal for Weight Sets 1 and 2, and with the winning side getting slightly better QoS for Weight Sets 3 and 4.

The preceding discussion has illustrated a qualitative approach to the choice of values for weights associated with a flow. At the present time, we have not performed extensive analysis or experiments regarding the options on how weight assignment affects resource allocation. Mathematically, there is no difference between applying a weight versus scaling all points on a utility curve. But from a practical point of view, the utility curve allows specification of quality attribute preferences (e.g., low vs. high resolution). Weights allow for indicating the relative importance of a user. Thus, weights and utility are two different questions. The experiments we described in this report represent a first look at these questions.

4 Experimental Results

4.1 Conducting the Experiment

During the experiment, we established and maintained all video flows as we discussed in Section 3.1.3. Below is a list of those flows.

- GVC from Left Squad to TOC
- GVC from Right Squad to TOC
- SC from Left Squad to TOC
- SC from Right Squad to TOC
- Local TOC stream (over high-capacity wired link)
- TOC stream to Left Squad
- TOC stream to Right Squad

All of these flows except the Local TOC stream were over a single, shared, wireless network. For the purpose of these experiments, we deliberately limited the bandwidth on the wireless link to 11Mbps in order to ensure that we were in a resource-constrained environment.

For each set of experimental conditions, we chose the inputs to the RMA so as to cause the desired squad to win the TOC camera and then collected video data over a three-minute interval. We recorded the number of received frames, the number of dropped frames, and the utility value for each received frame. We ran each experimental configuration a total of three times, and we used the average of the three runs in reporting our experimental data.

4.2 Resource Allocation Results

One of the goals of the experiment was to show how to use incentives to discourage overstatement of resource needs for allocation. In the experiments, the probability of being approached by a hostile was the only parameter that was (in some runs) overstated. Consequently, in the following description, overstatement refers to overstatement of that parameter.

4.2.1 Allocation Without Use of a Mechanism

We start with a baseline case in which no squad overstated the probability of hostile approach. Table 6 summarizes the input and results for this case. Each squad had an estimate of the probability of being approached by hostiles (shown as the *true* probability) and, in this case, reported that value without overstatement. The probabilities were simply low, medium, or high (L, M, or H). The *allocation* column shows which asset was allocated to each squad, with the GVC being the resource used by the squad that did not get the TOC camera. The *hostile approach* column has an *X* in the row corresponding to the squad that was approached by the hostile when the experiment was run. The *payment* column shows the payment the squad had to make; in the baseline case we used no mechanism to align incentives so we did not include this value (it would be equivalent to zero). The *value* is determined by the end state of the mission. In this case, thanks to having the TOC camera, the squad on the right (S_R) was able to detain the hostile that approached

it. The commander had defined that end state as having a value of 5, and that is the reward that squad S_R received. In this run of the experiment, the commander achieved the maximum possible mission value (i.e., 5). The *payoff* is the sum of the (negative) payment and the value for each squad. In cases summarized by Table 6 and Table 7, no mechanism was used, and consequently, there was no payment, and the payoff was equal to the value achieved.

Table 6: *Incentive Alignment Baseline (No Overstatement)*

Squad	Pr(Hostile) (True)	Pr(Hostile) (Report)	Allocation	Hostile Approach	Payment	Value	Payoff
Left	L	L	GVC		n/a	0	0
Right	M	M	TOC Camera	X	n/a	5	5
						Mission Value	5

Table 7 shows a case in which the squad on the left (S_L) overstates the probability of a hostile approach, perhaps with the intention of maximizing the chances of detaining the hostile, should he or she approach the left side. The allocation was made based on the reports to the RMA, with the result that the TOC camera was allocated to squad S_L . The hostile approached squad S_R , but since the GVC used by the squad deterred the hostile, he or she decided to abort the approach. As a consequence of not having the proper allocation of assets, no mission value was obtained. In particular, squad S_L (which overstated its resource need) did not undergo any negative consequence for overstating its resource needs. Nor did it contribute value to the overall mission.

Table 7: *Squad Left Overstatement Without Mechanism for Incentive Alignment*

Squad	Pr(Hostile) (True)	Pr(Hostile) Report	Allocation	Hostile Approach	Payment	Value	Payoff
Left	L	H	TOC Camera		n/a	0	0
Right	M	M	GVC	X	n/a	0	0
						Mission Value	0

Thus, the key point in comparing Table 6 and Table 7 is that when both squads accurately report their resource needs, the value of the mission is greater than when one squad overstates its resource needs. In contrast, when one of the squads overstates its resource needs, the overall mission value decreases. Since there is no consequence for overstatement, that behavior is not discouraged and so results in inefficient resource allocation.

4.2.2 Allocation that Includes a Mechanism

What is necessary, then, is a mechanism that will discourage possible overstatement of a resource request by a squad. The goal of incorporating a mechanism is to align the incentives of the two squads such that the overall mission value is maximized. The mechanism we used, dynamic VCG, uses payments to align incentives, as we demonstrate through the following examples.

Table 8 shows the same situation as Table 7, but with the mechanism in place to align incentives. In this case, the dynamic VCG computes payments by each squad for the resource it is allocated. These payments, shown in the table, represent the effect each squad has on others through its participation in the allocation (i.e., how much expected value a squad takes from others).¹⁵ Squad S_L must make a payment of 1.4; however, it does not earn any value because the hostile does not approach the left side. Consequently it has a negative payoff, that is, a loss.

Table 8: Squad Left Overstatement with Mechanism for Incentive Alignment

Squad	Pr(Hostile) (True)	Pr(Hostile) (Report)	Allocation	Hostile Approach	Payment	Value	Payoff
Left	L	H	TOC Camera		-1.4	0	-1.4
Right	M/	M	GVC	X	0	0	0
						Mission Value	0

With a mechanism in place, overstatements are discouraged. Table 9 shows the situation without overstatement. Squad S_R must make a payment for receiving allocation of the TOC camera. This payment is smaller than the one squad S_L made in the previous case. This is because the impact that squad S_R has on squad S_L (in terms of how much expected value is lost when squad S_L does not get the TOC camera) is lower. With the TOC camera, squad S_R is able to identify and detain the hostile that approached it. The squad receives a payoff of 4.4 and the commander achieves the maximum possible mission value, namely 5.

Table 9: No Overstatement with Mechanism for Incentive Alignment

Squad	Pr(Hostile) (True)	Pr(Hostile) (report)	Allocation	Hostile Approach	Payment	Value	Payoff
Left	L	L	GVC		0	0	0
Right	M	M	TOC Camera	X	-0.6	5	4.4
						Mission Value	5

In all the previous experiments, we assumed the hostiles approached the side that had the higher probability of a hostile approach. However, that will not always be the case because the side approached by the hostile is a stochastic event.¹⁶ For that reason, squad S_R may truthfully report the probability as M, get the TOC camera, and have to make a payment, but the hostile ends up approaching the other side. Squad S_R would not get any value, so it would finish with a negative payoff. Conversely, squad S_L could overstate the probability, get the TOC camera, and have the

¹⁵ This explains why the squads pay different amounts for the TOC camera, as evident in the differences between Table 8 and Table 9.

¹⁶ It is also possible for the hostiles to approach both sides of the base simultaneously. We do not explore this possibility because it will not change the outcome of how resources are allocated. That is, with one shared resource, if both sides are approached simultaneously, it will not affect the outcome.

hostile approaches the left side. In this case, squad S_L would finish with a positive payoff despite its overstatement. However, given the probabilities, such situations are less likely to occur than the situations depicted in the previous tables.¹⁷ In a single mission, a squad that overstates would be playing against the odds. On average, a squad achieves a higher payoff by being truthful.

To confirm this idea, we conducted a simulation in which we took into account all stochastic events, including hostile vs. friendly persons, the side the hostile intends to approach, the facial-recognition success, and whether or not hostiles are deterred by the GVC. Table 10 shows the result of simulating the scenario 10 times, with and without squad S_L overstating the probability. The numbers show that when squad S_L overstates its resource need, it receives on average a negative payoff, and average mission value is not as high as when squad S_L is truthful. When there is no overstatement of resource need, the commander achieves more average mission value, and the squads maximize their payoff. Therefore, squads' incentives are aligned with mission value.

Table 10: Simulation Results of Incentive Alignment

Squad Left Overstatement	Average Squad Left Payoff	Average Squad Right Payoff	Average Mission Value
yes	-1.2	0.5	1
no	0.5	3.6	4.5

The simulation allowed us to collect results for situations in which all the stochastic parameters had a uniform probability distribution, as opposed to the real experiments in which some of the stochastic parameters were kept fixed for practicality. The simulation confirmed that the mechanism was successful and maximized the expected mission value obtainable by the commander.

4.3 AQoS Results

To measure the effects of AQoS, we performed a series of experiments in which we varied the winner of the auction for allocation of the TOC camera and the weight set applied as a result of the outcome (Table 5). Weight Set 1, where weights of all flows are set to 1.0, served as a control case. We also performed an additional control case in which AQoS was disabled.¹⁸ A total of nine experimental configurations resulted.

We should note that due to the nature of the wireless radio environment and the adaptive nature of our approach, there is possible variation in results due to the dynamics of the environment. Thus, while it is not possible to *predict* a property such as bandwidth, our approach adjusts the allocation of bandwidth to respond to variation due to the environment.

In the following section, we will present our results in terms of the *unweighted* time-averaged utility τ . Unweighted τ values reflect the level of service that a particular *user* receives. Using weighted τ values would artificially confound weight and delivered service.

¹⁷ For the squad on the left, $\Pr(\text{Hostile} < 0.5)$. If the squad on the left overstates its resource need, it gets a positive payoff if hostiles approach it and a negative payoff if hostiles do not approach it. That is, $\Pr(\text{payoff} > 0) = \Pr(\text{Hostile}) < 1 - \Pr(\text{Hostile}) = \Pr(\text{payoff} < 0)$. Consequently, $\Pr(\text{payoff} > 0) < \Pr(\text{payoff} < 0)$.

¹⁸ When D-Q-RAM is disabled, there is no adaptation present, and all video flows attempt to operate at their highest QoS level.

4.3.1 TOC Camera

The output from the TOC camera is provided to the TOC as well as to either the left or right squad.

Figure 15 shows the experimental results for the allocation of the TOC camera in terms of time-averaged utility. For the TOC camera flows, we collected data at the TOC and at the left and right squads. The data collected at the TOC was transmitted over a high-bandwidth wired connection, so there were minimal dropped frames for this connection. The local TOC camera flow received the same weight as the winning squad's flow, to serve as a comparison.

As expected, for Weight Sets 2-4, the winning squad received high time-averaged utility (τ), and the losing squad received little or zero time-averaged utility (τ). We also notice that the value of τ received by the winning squad is nearly the same as the τ received locally (the bar labeled "TOC"), indicating that little degradation in quality occurred due to dropped frames over the wireless connection.

In the control case, defined by Weight Set 1, we see, as expected, that the left and right squads both receive about the same quality. We also see that the τ score each flow receives is lower than in the case where one squad is weighted higher than the other. This occurs because lowering the weight on one of the flows allows more resources to be applied to the flow for the squad that is allocated the TOC camera.

When D-Q-RAM is disabled, we see zero time-averaged utility (τ) for the left and right squads, and we see very high time-averaged utility (τ) for the local TOC flow. A closer examination of the data shows that the congestion that occurred without D-Q-RAM was so heavy that the squads were receiving no completed frames. The high τ for the local TOC stream, on the other hand, occurs because it is over a low-loss wired connection, and the sender is receiving no marginal utility messages to lower its operating point. As a result, it is operating at the highest possible QoS level with more than enough available bandwidth to support it at that level.




Figure 15: TOC Camera Allocation Time-Averaged Utility Scores

4.3.2 Ground Vehicle Camera

Figure 16 shows the experimental results for the GVC time-averaged utility. Since the GVC is intended as an alternative for the squad that does not have the TOC camera, the weights are set so that the loser receives better quality on that flow (in Weights Sets 2-4). We see from the GVC results that we get the expected outcome. When the right squad loses the auction for the TOC camera and controls the GVC, the τ score for the left (winning) squad is much lower than for the right squad. Similarly, when the left squad loses the TOC camera allocation, the τ score for the left squad is much higher than for the right squad.

On examining the control cases, we see that for Weight Set 1, both the left and right squads receive a high τ score. While this may seem acceptable, it in fact means that resources are being consumed for a task that does not greatly benefit from this consumption. We discussed this in Section 3.3 in relation to the TOC camera.

Finally, in the second control case, when D-Q-RAM is disabled, we see very poor performance for both the left and right squads. This result demonstrates the value of the D-Q-RAM approach.




Figure 16: Ground Vehicle Camera Time-Averaged Utility Scores

4.3.3 Squad Camera

Figure 17 shows the experimental results for time-averaged utility involving the squad camera. The set of QoS levels for the squad camera is identical to that for the GVC, but the weights are set differently in response to the allocation outcome. In the control case, with all weights equal, we get results comparable to those we achieved with the GVC. Unlike in the case of GVC allocation, we want some quality for both squads on the squad camera, regardless of the allocation outcome and use of weights that reflect this desire. In Weight Set 2, both left and right are given equal weight, which is reflected in the results shown in the figure. In Weight Sets 3 and 4, we give the winning side more weight, which is reflected in the slightly higher τ received by the winning side in our SC results.

When D-Q-RAM was disabled (as indicated in Figure 17), the results were somewhat unexpected. For the left squad we see that we measured a low τ score as expected, but for the right squad we unexpectedly measured a very high τ value. The most likely cause is that this flow received more bandwidth due to the location of the radio as compared to that of the TOC. Also, closer examination shows that even though the measured τ score is high, there is a 50% drop rate for this flow, suggesting that the τ score we measured for this case is somewhat misleading. The high score occurs because when D-Q-RAM is disabled, all frames that do get through are at the highest quality. As a result, even with a 50% loss rate, a reasonable τ score is still achievable.¹⁹

¹⁹ Under the current definition of τ , a flow in which all frames arrive with a utility of 0.5 will achieve the same τ as a flow in which only half the frames arrive with a utility of 1.0. This suggests that it may be worthwhile to consider a revision of the definition of τ , to more heavily penalize the dropping of frames.




Figure 17: Squad Camera Time-Averaged Utility Scores

4.3.4 Dropped Frames

In addition to measuring the τ score for each flow, we also measured the fraction of frames that were dropped. As an example, we show the results for the left squad camera in Figure 18. Results for other flows are similar. We find that for all the cases where we enable D-Q-RAM, the fraction of dropped frames is very low, regardless of the weights we apply to the flows.²⁰ Note that disabling D-Q-RAM resulted in very high loss, over 90% for most of the flows we recorded. These measured results are supported by our qualitative observations from the experiment. We found that, visually, the video produced with D-Q-RAM enabled was smooth and uninterrupted, but it was very choppy and contained frequent long pauses when D-Q-RAM was disabled.

²⁰ When a frame is sent, it is composed of multiple packets. If a packet is lost, the entire frame is considered lost.




Figure 18: Fraction of Frames Dropped From Left Squad Camera

4.4 Summary of Experimental Results

We offer three points as a summary to the experimental results provided in this section. First, the use of computational mechanism design, more specifically dynamic VCG, would discourage a rational agent from misreporting needs, allowing the aggregation of truthful information needed to allocate resources optimally. Second, AQoS demonstrates effective management of network bandwidth in the presence of uncertainty due to dynamic effects. Finally, we have consistently applied these approaches to address the need for resource allocation and adaptation.

5 Summary

This report has described our scenario-based approach to examining resource allocation when a request for a resource involves a possible overstatement of resource needs. The scenario involves small groups of soldiers who are using a shared camera. We examined the case in which one group overstates a resource need, thereby denying others the use of that resource, which ultimately may affect mission success. In an environment subject to dynamic effects, it is necessary to provide a means to adapt to network characteristics, notably bandwidth. We have addressed this aspect of the problem by taking an AQoS approach.

The key results of this work are summarized as follows:

- The use of a computational mechanism design approach would discourage rational agents from overstating resource needs, allowing the aggregation of truthful information needed to allocate resources optimally.
- An AQoS approach has demonstrated optimality in allocating network bandwidth in the presence of unpredictable resource availability. This result can apply to many different settings; for example, to a wireless network due to its inherently unpredictable nature.
- The computational mechanism design and AQoS approaches have been used to allocate available network bandwidth consistent with sensor allocation.

Although the above results demonstrate successful use of computational mechanism design and AQoS, additional issues must be resolved to establish the operational viability of this approach. Regarding the use of computational mechanism design, first, there is a need to include mechanisms that do not require the use of a payment to align incentives. Second, the current approach allocates a resource using dynamic VCG, but it does not account for a possible overstatement in the utility values associated with that resource. Third, the various methods by which the allocation decision is made should be examined for their practical implications (e.g., allocation at a periodic interval, upon a report from a user, or upon a change in the network state such as detection of a fault). Fourth, it is imperative that an approach be developed that is acceptable to a user and that hides the details of the underlying mathematics; we made a start in that direction in this report. Regarding the use of AQoS, we applied the current approach in a single-hop network, which must be scaled to address the more general multi-hop case. In addition, a thorough investigation of the assignment and use of weights in regard to resource allocation is worthy of further investigation. Finally, there is a need to address an integrated approach to assess incentive compatibility and its optimality character.

Acronym List

AQoS	Adaptive Quality of Service
CMD	computational mechanism design
D-Q-RAM	Distributed QoS Resource Allocation Model
FOB	Forward Operations Base
FRS	Facial-Recognition Server
GCS	Ground Control Station
GVC	Ground Vehicle Camera
ISR	Intelligence, Surveillance and Reconnaissance
MDP	Markov Decision Process
Q-RAM	Quality of Service Resource Allocation Model
RMA	Resource Management Agent
TAU	time-averaged utility
TCP	Transmission Control Protocol
TOC	Tactical Operations Center
UAV	unmanned aerial vehicle
UDP	User Diagram Protocol
VCG	Vickrey-Clarke-Groves

Appendix A: Scenario Details

This appendix provides a step-by-step description of the scenario that we presented in the document. Here we introduce two cases. The first involves an overstatement of a resource request (Section A.1), and the second is a modification of that case involving no overstatement of a resource request (Section A.2).

A.1: Case of Resource Overstatement

Table 11 describes the case in which there is an overstatement for a resource request (the tower camera).

Table 11: Details of Scenario Steps

Topic	With Misreporting
Pre-deployment	<ul style="list-style-type: none">• Setpoints have been determined for images based on values of resolution and compression.• Utility curves have been developed for video images.• Weights have been determined for each resource based on mission and expected outcome, but the resource has not been assigned to a particular squad. This also applies to embedded resources.• The commander and the squad leaders have met and determined initial positions for the TOC and the two squads, as well as other, related procedural matters of relevance.• Rewards and value have been determined by the commander.• The FRS has been pre-loaded with images of hostiles.
Initial deployment	<ul style="list-style-type: none">• Squads S_L and S_R are deployed to their pre-planned positions. No shared resources, such as the TOC camera, have been allocated to the squads.• Notification is sent to the TOC that the squad cams are available.
TOC access to squad cams	The TOC requests that images from each squad camera are sent to the TOC. They are sent and displayed at the TOC.
TOC camera report	After the TOC camera is deployed and is operational, it reports to the RMA that it is in a <i>ready</i> state (i.e., it may now be assigned by the RMA via CMD).
GVC cameras report	When the two GVC cameras are initially operational, they report to the TOC that they are in a <i>ready</i> state.
RMA “announcement”	After some time period following the initial report that some device is ready (in this case the TOC camera), the RMA reports to the users (the two squads) that a resource is available.
S_L sends preferences to RMA	<ul style="list-style-type: none">• S_L is deployed and aware that resources are available for allocation. It now sends its preferences to the RMA.• In its report S_L overstates the probability of hostiles approaching its side. It expects that the probability will be <i>Low</i> but purposefully reports a value of <i>High</i>. It does this because it can get a higher reward if it captures the hostiles.
S_R sends preferences to RMA	<ul style="list-style-type: none">• This step is similar to the step above.• S_R reports probabilities that are assumed to be correct (i.e., not overstated). It reports a value of <i>Medium</i> for the probability of a hostile approaching its side.

Table 11: Details of Scenario Steps (cont'd.)

Topic	With Misreporting
RMA response to S_L and S_R	The resource manager agent, using a CMD approach, executes its assignment algorithm based on the preferences that S_L and S_R reported. The results are <ul style="list-style-type: none"> S_L receives the TOC camera. S_R is expected to use its GVC.
TOC camera control by S_L	The TOC camera is now under control of S_L . It steers the camera such that it can obtain a view of the expected approach corridor where hostiles may come.
GVC camera control by S_R	S_R uses the GVC.
S_L identifies unknowns	<ul style="list-style-type: none"> A number of unknowns approach S_L in a vehicle. S_L stops the vehicle and performs facial recognition of the occupants by sending images to the Facial-Recognition Server (FRS). There are no facial-recognition matches indicating an occupant is a hostile, and the individuals are allowed to pass.
S_R identifies unknowns	<p>A number of individuals approach S_R in a vehicle.</p> <p>S_R stops the vehicle and performs facial recognition of the occupants by sending images to the FRS.</p> <p>There are no facial-recognition matches indicating an occupant is a hostile, and the individuals are allowed to pass.</p>
S_L observations	S_L , using the TOC camera, notices a pickup truck approaching with about eight people in it. This appears to be an unusual situation and it is reported to the TOC.
TOC requests images	<ul style="list-style-type: none"> Due to ongoing activities at S_L, the TOC also requests images from the TOC camera and the GVC, to have better situational awareness. The TOC also requests that images from the S_L squad camera be provided with a higher resolution over a broader area.
Special person	A special person who may be able to identify a hostile is at the TOC. The person requests that high-resolution video images from the handheld facial-recognition devices be sent and displayed at the TOC. Further details are unavailable due to security concerns.
Network overload	The combination of sending images to both the TOC as well as the FRS overloads the network.
Network response	AQoS is invoked. Network transfers associated with the higher priority mission (and thus assumed higher priority QRAM weights) are affected and assigned weights that will affect network transmission.
S_L facial recognition complete	S_L obtains responses from the FRS indicating that the individuals are not in a database of known hostiles. The individuals are allowed to pass.
S_R facial recognition complete	<ul style="list-style-type: none"> S_R obtains facial-recognition results from the FRS server. There are no facial-recognition matches indicating a vehicle occupant is a hostile, and the individuals are allowed to pass.
Hostiles approach S_R	Several unknowns approach S_R in what appears to be a normal manner, but they are actually hostiles.
Hostiles escape	The hostiles detect a GVC camera on the road and are deterred from continued approach to S_R , fearing that they may be identified and captured. They change their path so that they are not intercepted by squad S_R and escape.

A.2: Non-Overstatement of Resource Request

The scenario steps in Table 11 have illustrated the case where an overstatement of resource need resulted in a failure to detect hostiles. In the absence of an overstatement of resource need by S_L , the TOC camera would have been allocated to S_R instead of S_L . This would have allowed S_R to better identify and possibly to capture the hostiles that approached in their direction.

Appendix B: Utility Curves Used in the Scenario

In this appendix, we provide utility curves for the TOC camera, Ground Vehicle Camera, squad camera, and smartphone (for facial-recognition images) that were used in the scenario. The utility curves are associated with distribution of images from a type of camera.

TOC Camera

For the TOC camera, we illustrate the choice of setpoints used to construct the utility function. The relevant values of utility were *pairs* of values based on the camera resolution and compression. We determined that the user wanted a frame rate of 15 frames per second. For a given pair (resolution and compression²¹) we experimentally determined the value of the bandwidth (in kilobits per second) for a particular camera. We chose the value of bandwidth as the setpoint and a user provided the value of utility. The resulting TOC camera utility values and marginal utility values are shown in Table 12. We determined utility values for other video devices in a similar manner and the results are shown below.

Table 12: Sample Utility Values for TOC Camera

Resolution	Compression	Bandwidth (kbit/s)	Utility	Marginal Utility
88x60	30	204.7208	0	7.44×10^{-3}
88x60	40	222.404	0.1315	5.09×10^{-3}
88x60	50	243.9376	0.241138	4.68×10^{-3}
88x60	55	254.1472	0.288919	3.90×10^{-3}
88x60	60	265.3352	0.332549	3.45×10^{-3}
88x60	65	276.8856	0.372387	1.81×10^{-3}
88x60	70	296.9936	0.408763	1.80×10^{-3}
88x60	75	315.4808	0.441978	1.14×10^{-3}
88x60	80	342.1336	0.472307	6.96×10^{-4}
88x60	85	381.9672	0.5	3.82×10^{-4}
176x120	75	817.136	0.666198	3.93×10^{-4}
176x120	80	910.232	0.696527	1.45×10^{-4}
176x120	85	1097.064	0.72422	7.01×10^{-5}

²¹ The value of compression ranged from 0 to 100 with a larger value denoting a higher quality image, that is, less compression. The value of resolution is expressed in pixels.

Table 12: Sample Utility Values for TOC Camera (con't.)

Resolution	Compression	Bandwidth (kbit/s)	Utility	Marginal Utility
352x240	80	2992.712	0.857188	4.55×10^{-5}
352x240	85	3600.12	0.884881	2.97×10^{-5}
704x480	85	7465.528	1.0	

Figure 19 shows a plot of the utility curve for the TOC camera, and all of the points are on the concave majorant.




Figure 19: Utility Curve for the TOC Camera

It is interesting to note that the bandwidth, chosen as the setpoint, was a function of *two* other parameters, namely resolution and compression, for a fixed frame rate. This choice illustrates the versatility of the AQoS.

Ground Vehicle Camera and Squad Cameras

Each squad may use a GVC as well as a squad camera. A squad uses the GVC when it is not provided access to the TOC camera. Each squad uses the squad camera to provide situational awareness to the commander in the TOC.

Figure 20 shows the utility curve that was used for both the GVC and the squad camera.




Figure 20: Utility Curves for Ground Vehicle Camera and Squad Camera

Facial Recognition

Facial recognition was used as a means to determine if an individual was a hostile. Using a smartphone, a squad agent captured an image of the individual and sent it to the FRS. The FRS would then respond as to whether the image matched a known hostile.

Figure 21 shows the utility curve for the smartphone camera used in facial recognition.




Figure 21: Facial-Recognition Utility Values

Summary

Notice that the utility curves presented above are similar in shape and that they satisfy the conditions necessary for achieving optimality in the AQoS approach. In addition, a given utility curve also manifests characteristics of the camera with which it is associated. For example, the TOC camera could require greater bandwidth, due to the greater values of resolution and image compression available. This is also true for the use of the smartphone camera, where such values were smaller than for other cameras; hence, its bandwidth requirements were correspondingly smaller. Thus, the maximum bandwidth consumed by the smartphone was 250 kBits/sec, while for the TOC camera it was over 7000 kBits/sec.

References

URLs are valid as of the publication date of this document.

[Bergemann 2006]

Bergemann, D. & Valimaki, J. “Efficient Dynamic Auctions,” *Cowles Foundation Discussion*, Paper No. 1584, 2006. <http://cowles.econ.yale.edu/P/cd/d15b/d1584.pdf>

[Cavallo 2009]

Cavallo, R; Parkes, D. C.; & Singh, S. *Efficient Mechanisms with Dynamic Populations and Dynamic Types*. Harvard University, 2009.

[Clarke 1971]

Clarke, E. “Multipart Pricing of Public Goods,” *Public Choice* 2 (1971):19-33.

[Dash 2003]

Dash, R.K.; Jennings, N.R.; & Parkes, D.C.; “Computational-Mechanism Design: A Call to Arms.” *IEEE Intelligent Systems* 18, 6 (Nov-Dec 2003): 40-47. IEEE, 2003.

[Greenberg 2010]

Greenberg, K. “Unmanned Aerial Systems: Quality As Well As Quantity.” *Military Review* 90, 4 (July-August, 2010): 51-59.

[Groves 1973]

Groves, T. “Incentives in Teams.” *Econometrica*, 41 (1973): 617-631.

[Hansen 2010]

Hansen, J.; Hissam, S.; Meyers, B. C.; Morris, E.; Plakosh, D.; Simanta, S.; & Wrage, L. *Adaptive Flow Control for Enabling Quality of Service in Tactical Ad Hoc Wireless Networks*, (TR-2010 - 030). Software Engineering Institute, December 2010.
<http://www.sei.cmu.edu/library/abstracts/reports/10tr030.cfm>

[Hansen 2012]

Hansen, J.; Hissam S.; Plakosh, D.; & Wrage, L “Adaptive Quality of Service in Ad Hoc Wireless Networks.” *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC 2012)*. Paris, France, April 2012. IEEE, 2012.

[Hoover 2001]

Hoover, C.; Hansen, J.; Koopman, P.; & Tamboli, S. “The Amaranth Framework: Policy-Based Quality of Service Management for High Assurance Computing.” *International Journal of Reliability, Quality and Safety Engineering* 8, 4 (December 2001): 323-350.

[Kuhn 1951]

Kuhn, H.W. & Tucker, A.W. “Nonlinear Programming” (482-492). *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley, CA, July-August 1950. University of California Press, 1951.

[Lee 1999]

Lee, Chen; Lehoczky, John; Rajkumar, Raj; & Hansen, Jeffery P. “A Scalable Solution to the Multi-Resource QoS Problem” (315-326). *Proceedings of the 20th IEEE Real-Time Systems Symposium*, Phoenix, AZ, December 1999. IEEE, 1999.

[US Army 1991]

United States Army. *Field Manual 24-17: Tactical Record Traffic System*, Department of the Army, 1991.

[Vickrey 1961]

Vickrey, W. “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *The Journal of Finance* 16, 1 (March 1961): 8-37.

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE October 2012		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Resource Allocation in Dynamic Environments			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Jeff Hansen, Scott Hissam, B. Craig Meyers, Gabriel Moreno, Daniel M. Plakosh, Joseph Seibel, Lutz Wrage				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) SEI Administrative Agent ESC/CAA 20 Schilling Circle, Building 1305, 3rd Floor Hanscom AFB, MA 01731-2125			10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2012-011	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) This technical report examines two challenges related to resource allocation that can negatively affect system operation in a dynamic environment, where warfighter needs for resources, resource availability, environmental effects, and mission conditions can change from moment to moment. The first challenge occurs when warfighters overstate their individual needs of a shared resource, leading to inefficient allocation. Overstatement may bring local optimization; however, it can cause global inefficiencies that result in a detriment to overall mission success. This challenge is addressed by using computational mechanism design, more specifically, the dynamic Vickrey-Clark-Groves allocation mechanism. The second challenge involves resource availability that may change frequently. Such is the case in a wireless mesh network where routes and bandwidth may vary over even small intervals of time. In such a case, an adaptive quality of service (AQoS) approach is used, and the available resource is allocated using the Dynamic QoS-based Resource Allocation Model (D-Q-RAM). Computational mechanism design is used to allocate sensors, and the AQoS approach allocates the available network bandwidth in a way consistent with the sensor allocation, providing an approach for dealing with resource allocation and adaptation in a dynamic environment. Initial experimental results of applying the approach are reported.				
14. SUBJECT TERMS Resource allocation, Adaptive Quality of Service			15. NUMBER OF PAGES 59	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	